arXiv:1602.05765v1 [cs.AI] 18 Feb 2016

Entity Embeddings with Conceptual Subspaces as a Basis for Plausible Reasoning

Shoaib Jameel
Cardiff University
United Kingdom
JameelS1@cardiff.ac.uk

Steven Schockaert
Cardiff University
United Kingdom
SchockaertS1@cardiff.ac.uk

Abstract

Conceptual spaces are geometric representations of conceptual knowledge, in which entities correspond to points, natural properties correspond to convex regions, and the dimensions of the space correspond to salient features. While conceptual spaces enable elegant models of various cognitive phenomena, the lack of automated methods for constructing such representations have so far limited their application in artificial intelligence. To address this issue, we propose a method which learns a vector-space embedding of entities from Wikipedia and constrains this embedding such that entities of the same semantic type are located in some lower-dimensional subspace. We experimentally demonstrate the usefulness of these subspaces as (approximate) conceptual space representations by showing, among others, that important features can be modelled as directions and that natural properties tend to correspond to convex regions.

1 Introduction

Despite the fact that several large-scale open-domain knowledge bases are now available (e.g. CYC, SUMO, Freebase, Wikidata and YAGO), few knowledge-driven applications rely on logical reasoning. An important reason for this is that available knowledge is often inconsistent. For example, the concept *ice cream shop* is asserted to be disjoint from *restaurant* in CYC, while it is considered a type of *restaurant* on Wikipedia[1]. Another challenge for logical reasoning is that available knowledge is seldom complete. For example, SUMO encodes[2] knowledge about chess, darts and poker, but mentions nothing about checkers.

Humans are remarkably adept at overcoming such challenges [Collins and Michalski, 1989; Festinger, 1957]. For example, we can recognize that the aforementioned conflict between CYC and Freebase is caused by the vagueness of the categories *restaurant* and *shop*, which both have *ice cream shop* as a borderline case. Similarly, we can deal with knowledge gaps by making inductive inferences, e.g. assuming that properties which hold for chess, darts and poker should hold for checkers as well. Automating such forms of plausible reasoning has proven challenging, among others because they rely on an underlying notion of similarity, which is difficult to characterize using purely symbolic methods.

The solution offered by the theory of conceptual spaces [Gärdenfors, 2000] is to represent concepts as regions in a suitable metric space. The points of this space correspond to (actual or possible) entities of a given semantic type, such that similar entities are located close to each other. It is furthermore posited that most natural properties will correspond to regions that are convex, in accordance with prototype theory [Rosch, 1973]. Furthermore, the dimensions of a conceptual space correspond to the salient features of the considered domain. For example, a conceptual space of wines could have dimensions relating to sweetness, acidity, fruitiness, amount of tannins, etc. Using conceptual space representations, many cognitive phenomena, including vagueness and induction, can be modelled in a natural way [Gärdenfors, 2000; Douven *et al.*, 2013; Schockaert and Prade, 2013; Lieto *et al.*, 2015]. However, existing applications have focused on a few particular domains where conceptual space representations can be derived from available metric information. For example, several authors have considered conceptual spaces for music perception [Forth *et al.*, 2010; Chella, 2015]. In such cases, the definition of the conceptual space, and its relationship to e.g. audio signals, relies on well-understood insights from the field of music cognition.

The research question we consider in this paper is whether we can automatically obtain approximate conceptual space representations for a wide range of domains, by combining information found in existing knowledge bases with representations derived from large text corpora such as Wikipedia.

Our approach builds on existing work for learning word embeddings from text corpora. Word embeddings [Mikolov *et al.*, 2013; Pennington *et al.*, 2014; Turney and Pantel, 2010] represent the meaning of words as points in a high-dimensional Euclidean space, and are in this sense reminiscent of conceptual spaces. However, they differ from conceptual spaces in at least two crucial ways, which limits their usefulness for applications in knowledge representation (e.g. automatically repairing inconsistencies). First, because conceptual spaces represent properties and

[1] https://en.wikipedia.org/wiki/Category:Types_of_restaurant
[2] https://github.com/ontologyportal/sumo/blob/master/Merge.kif

concepts as regions, semantic relations such as subsumption, overlap and typicality can be naturally modelled. For example, we can encode that an ice cream shop is an atypical kind of shop which is similar to a restaurant by representing ice cream shop as a region which is included in the region for shop and disjoint from but close to the region for restaurant. In contrast, most models for word embedding are only aimed at modelling similarity (and related notions such as analogy). Second, the dimensions of a conceptual space directly reflect the salient properties of the underlying domain, which allows us to compare or rank entities, to model context effects[3], and to describe how two entities or concepts are semantically related (e.g. that the rules of chess are more complex than the rules of checkers). In contrast, the dimensions of a word embedding space are essentially meaningless, which is a consequence of the fact that word embeddings are not tied to a specific domain.

The solution proposed in [Derrac and Schockaert, 2015] is to learn a different vector space representation for each semantic type (e.g. movies), given a textual description of the entities in that domain (e.g. movie reviews). Specifically, they use multi-dimensional scaling (MDS) to construct the space and identify directions corresponding to salient properties of the considered domain in a post-hoc analysis. The model we propose improves this approach in several ways. First, we learn a single vector space, which has a particular subspace for each semantic type. Among others, this allows us to model semantic type hierarchies and to take into account relations between entities of different semantic types to align their corresponding subspaces (e.g. the subspaces representing actors, directors and genres can help to obtain a more accurate representation of movies). Second, our model is specifically aimed at finding a representation in which the salient properties of a given domain correspond to directions in the corresponding subspace. Third, MDS requires a distance matrix whose size is quadratic in the number of entities, which severely limits the scalability of the method from [Derrac and Schockaert, 2015]. In contrast, our model can easily learn representations for millions of entities.

2 Related work

Word embedding Word embeddings are vector space representations which are used to model the meaning of words. Several existing models construct a vector for each word by applying some form of matrix factorization to a term-term co-occurrence matrix; see [Turney and Pantel, 2010] for an overview of such approaches. Recently, a number of models have been proposed which instead explicitly optimize the predictive power of the word vectors. For example, the popular Skip-gram model [Mikolov *et al.*, 2013] tries to find word vectors that can be used to predict the probability of seeing a context word, given an occurrence of the word being modelled, while the related continuous bag-or-words (CBOW) model focuses on the probability of seeing the word being modelled, given the occurrence of a context word.

[3]The context-dependent nature of similarity is modelled in conceptual spaces by allowing dimensions to be rescaled, depending on the importance of the corresponding property in the given context.

An interesting property of word embeddings is that they often capture several kinds of semantic relations, beyond simple similarity. For example, in [Mikolov *et al.*, 2013] it is shown that analogical proportions of the form a is to b what c is to d correspond to approximate parallelograms in the space obtained by Skip-gram. They also found that vector addition sometimes corresponds to a form of semantic composition, e.g. adding the vectors for *Germany* and *capital* resulted in a vector which is close to the vector for *Berlin*.

The fact that the vector space obtained by the Skip-gram model satisfies such linear regularities is at first glance somewhat surprising. In [Pennington *et al.*, 2014], the authors analyze what characteristics of a word embedding model can explain this effect, and propose a new model, called GloVe, which is explicitly aimed at capturing linear regularities. Since our model will build on GloVe, we briefly review its formulation. The GloVe model relies on a term-term co-occurrence matrix $X = (x_{ij})$, where x_{ij} is the number of times that word i appears in the context of word j. For each term t_i in the vocabulary, two word vectors w_i and $\tilde{w}_i$ and a bias b_i are chosen by minimizing the following objective:

$$J = \sum_{i=1}^{V} \sum_{j=1}^{V} f(x_{ij})(w_i \cdot \tilde{w}_j + b_i + b_j - \log x_{ij})^2 \qquad (1)$$

where V is the number of words in the vocabulary. The function f is used to prevent common words from dominating the objective function, and is defined as follows:

$$f(x_{ij}) = \begin{cases} \left(\frac{x_{ij}}{x_{max}}\right)^{\alpha} & \text{if } x_{ij} < x_{\max} \\ 1 & \text{otherwise} \end{cases} \qquad (2)$$

where x_{max} is a constant which was fixed as 100. Intuitively, w_i reflects the meaning of term t_i while $\tilde{w}_i$ reflects how the occurence of that term in the context of another term t_j impacts the meaning of t_j.

Knowledge graph embedding Knowledge bases such as Freebase and Wikidata can essentially be seen as collections of (subject, predicate, object) triples, and can thus be encoded as a graph, where nodes correspond to entities and edges are labelled with relation types. Several authors have looked at the problem of automatically expanding such knowledge graphs [Dong *et al.*, 2014]. Here, we focus on models that rely on embedding knowledge graphs in a vector space, as we will use similar ideas for aligning different conceptual subspaces. The idea of embedding knowledge graphs in a vector space was proposed in [Bordes *et al.*, 2011]. In particular, they propose the model SE, in which each entity e_i is represented as a vector and each relation r_k is represented using two matrices R_k^{lhs} and R_k^{rhs}. The constraint they impose is that the following distance should be small for triples (e_i, r_k, e_j) in the knowledge graph and large for other triples:

$$d(R_k^{lhs} e_i, R_k^{rhs} e_j)$$

where d is either the Euclidean or Manhattan distance. An important drawback of this model is that it requires learning a large number of parameters, which leads to underfitting. In

[Bordes *et al.*, 2013] a simpler alternative, called TransE, was proposed, which represents reach relation as a vector and considers the following scoring function instead:

$$d(e_i + r_k, e_j)$$

Despite the simplicity of this model, it was shown to substantially outperform SE in practice. However, as noted in [Wang *et al.*, 2014], TransE is mostly suitable for one-to-one relations. To obtain a more faithful modelling of one-to-many, many-to-one and many-to-many relations, the model TransH is proposed. In this model, both a hyperplane H_k and an $(n-1)$ dimensional vector r_k is associated with each relation type (with n the dimension of the embedding space), and the following scoring function is considered:

$$d(e_i^{H_k} + r_k, e_j^{H_k})$$

where $e_i^{H_k}$ and $e_j^{H_k}$ are the orthogonal projections of e_i and e_j on the hyperplane H_k. The TransR model, introduced in [Lin *et al.*, 2015], follows a similar strategy, but instead associates an m-dimensional vector r_k and an $m \times n$ matrix M_k with each relation, and uses the following scoring function:

$$d(e_i M_r + r_k, e_j M_r)$$

The underlying idea is to use the TransE model, after projecting the entities onto a relation-specific space. While in general it is not required that $n = m$, this particular choice was used in all experiments. Finally, [Lin *et al.*, 2015] also proposes a variant CTransR, in which each entities are clustered, and each relation can have a different representation for each cluster.

In our model, the semantic types of entities play a crucial role. One other approach that explicitly takes semantic type into account is [Guo *et al.*, 2015]. In particular, they add a regularization term to the objective function of existing embedding models to encode the requirement that entities of the same semantic type should be represented using similar vectors, which they formalize based on two manifold learning algorithms. Unfortunately, the scalability of the resulting method is relatively limited.

In [Wang *et al.*, 2014] a model is proposed that combines word embedding with knowledge graph embedding. In particular, they jointly learn a representation for words, entities and relations, where the word representations are constrained similarly as in the Skip-gram model and the entities and relations are constrained similarly as in the TransE model. The entity and word representations are aligned either based on Wikipedia anchors or based on the entity names. An improvement of this method was proposed in [Zhong *et al.*, 2015], where the alignment is instead based on the text of the Wikipedia article of the entity. Along similar lines, [Xu *et al.*, 2014] proposes a model in which the objective functions of Skip-gram and TransE are combined. A third component in their objective function allows the model to take into account an external similarity relation, by imposing the requirement that similar terms should have similar vectors. It is shown that the resulting model improves the word embeddings from Skip-gram. However, since this model requires us to identify entities with their name, its usefulness for knowledge graph embedding is limited.

3 Description of the model

Our aim is to learn a vector-space embedding of a set of entities E, in which entities of the same semantic type lie in some lower-dimensional subspace. Let S be the set of all semantic types. For $s \in S$, we write E_s for the set of all entities of type s. We furthermore assume that a set of binary relations R is available, and a set $G \subseteq E \times R \times E$ of triples of the form (e, k, f), encoding that entities e and f are in relation k. Finally, we assume that for every entity e, a bag of words W_e describing that entity is available. The model we propose has the following form:

$$J = \alpha J_{text} + (1-\alpha)(J_{type} + J_{rel}) + \beta J_{reg} \quad (3)$$

where $\alpha \in [0,1]$ and $\beta \in [0, +\infty[$ are parameters controlling the relative importance of the different components of the model. Component J_{text} will be used to constrain the representation of the entities based on their textual description, J_{type} will impose the constraint that entities of the same type belong to a particular subspace, J_{rel} will use the relations in R to improve the alignment between these subspaces, and J_{reg} is a regularization component which will allow the model to automatically select the most appropriate number of dimensions for every subspace. We now discuss each of these components in more detail.

Entity embedding From the bag of words representations W_e, we want to find a point $p_e \in \mathbb{R}^n$ for each entity e such that similar entities correspond to nearby points and such that salient features can be interpreted as directions in the space. Specifically, let f be a feature of interest, and let $x_i \in \mathbb{R}$ be the value of feature f for entity e_i, i.e. x_i reflects how much e_i has feature f. Then there should be a vector $w_f \in \mathbb{R}^n$ such that the orthogonal projection p'_{e_i} of the point p_{e_i} on the line $L_f = \{q \,|\, q = \lambda \cdot w_f, \lambda \in \mathbb{R}\}$ is given by $p'_{e_i} = c_f x_i w_f$ for some constant $c_f \in \mathbb{R}$. In other words, in a coordinate system where L_f coincides with one of the axes, the corresponding coordinate of p_e should be proportional to x_i. This requirement can be imposed as follows[4]:

$$p_{e_i} \cdot w_f \propto x_i \quad (4)$$

Unfortunately, we do not actually know what are the salient features in most domains. Following [Derrac and Schockaert, 2015], we therefore use word co-occurrence as a proxy for feature values. In particular, we assume that each word potentially corresponds to a salient feature, and that the number of times a word co-occurs with a given entity reflects how much that entity has the corresponding feature. This leads to the following constraint

$$p_{e_i} \cdot w_j = g(y_{ji}) \quad (5)$$

where y_{ji} is the number of times word t_j occurs in W_{e_i} and g is a monotonic function that maps co-occurrence statistics to feature values. Typically it will not be possible to satisfy the constraint (5) for all entities and all context words. The assumption underlying this model is that the salient features

[4]We are abusing notation here, using p_{e_i} as a notation for the vector $\overrightarrow{0p_{e_i}}$.

of an entity affect the co-occurrence statistics of many context words, and that the words for which (5) is (approximately) satisfied, in an optimal solution, will therefore be those that are strongly related to important features of the entity e_i.

Note that the requirement in (5) closely resembles the constraints that are optimized by the GloVe model. Moreover, as in the GloVe model, we can choose $g(y_{ji}) = \log(y_{ji})$ and formalize the objective function as a least squares regression problem, weighted such that frequent terms have a stronger impact on the objective function:

$$J^E_{\textit{text}} = \sum_{e_i \in E} \sum_{t_j \in W_{e_i}} f(y_{ji})(p_{e_i} \cdot w_j + b_i + b_j - \log y_{ji})^2$$

where f is defined as in (2). The resulting model is essentially the same as GloVe, but instead of modelling word-word co-occurrence we now model entity-word co-occurrence. The geometric interpretation, however, is different, as we view entities as points and context words as vectors. We can further constrain the word vectors w_j by adding a second component which corresponds to the original GloVe model. In particular, we define $J_{\textit{text}} = J^E_{\textit{text}} + J_{\textit{glove}}$, where $J_{\textit{glove}}$ is the objective function J defined in (1).

Subspace constraints To impose the constraint that all entities of a given type s should belong to the same subspace, we associate with each semantic type s a set of $n+1$ points $p^s_0, ..., p^s_n \in \mathbb{R}^n$ and express that for each entity e_i of type s, the point p_{e_i} can be written as a linear combination of the points $p^s_0, ..., p^s_n$:

$$J_{\textit{type}} = \sum_{s \in S} \sum_{e \in E_s} \Big(p_e - \sum_{j=0}^{n} \lambda^{e,s}_j p^s_j \Big)^2$$

Note that on its own, this component is trivial, as it suffices to choose any set of points $p^s_0, ..., p^s_n$ in general linear position. However, we will additionally require that the space spanned by the points $p^s_0, ..., p^s_n$ is as low-dimensional as possible. In particular, let M_s be the $n \times n$ matrix whose i^{th} row vector is $p^s_i - p^s_0$. Then clearly the rank of M_s is equal to the dimension of the space spanned by $p^s_0, ..., p^s_n$. We now want to add a regularization term to penalize high-rank matrices M_s. Unfortunately, no efficient methods exist for directly minimizing the rank of a matrix M. The relaxation suggested in [Fazel, 2002] is to minimize the nuclear norm $\|M\|_*$ instead (i.e. the sum of the singular values of M). This technique was empirically shown to lead to low-rank matrix solutions in many applications, and is known to be equivalent to rank minimization in certain cases [Recht *et al.*, 2010]. The regularization term associated with $J_{\textit{type}}$ is thus given by

$$J^1_{\textit{reg}} = \sum_{s \in S} \|M_s\|_*$$

To implement nuclear norm regularization, we have used the recently proposed method from [Hsieh and Olsen, 2014].

Modelling relations Often we have information about how entities of different types are related, e.g. the fact that Steven Spielberg is the director of Jurassic Park. Such relationships can help us to align the subspaces corresponding to different types. Since our main aim is to improve the entity embeddings, rather than predicting relationships between entities of different types, methods such as TransH and TransR, which rely on projecting the entities to a different space, are not directly suitable. On the other hand, TransE is only suitable for one-to-one relations.

We propose an alternative to TransE which is inspired by our modeling of semantic types. As in TransE, we assume that every relation k is represented as a vector r_k. We furthermore write $\textit{rhs}(e,k) = \{f \,|\, (e,k,f) \in G\}$ and $\textit{lhs}(k,f) = \{e \,|\, (e,k,f) \in G\}$. Rather than imposing that $e + r_k = f$ if $(e,k,f) \in G$, as in TransE, we require that the points in $P_{e,k} = \{p_f \,|\, f \in \textit{rhs}(e,k)\} \cup \{p_e + r_k\}$ lie in a low-dimensional subspace and, similarly, that the points in $P_{k,f} = \{p_e \,|\, e \in \textit{lhs}(k,f)\} \cup \{p_f - r_k\}$ lie in a low-dimensional subspace. Note that in the case of one-to-many or many-to-one relations, this part of the model is similar to TransH in the special case where the considered subspaces are one-dimensional. Similar to TransE, we additionally impose that the points for the entities in $\textit{rhs}(e,k)$ are all close to $p_e + r_k$ and that the points for the entities in $\textit{lhs}(k,f)$ are all close to $p_f - r_k$. The resulting objective $J_{\textit{rel}}$ is given by:

$$\sum_{k \in R} \sum_{p \in P_{e,k}} \big(p - \sum_{j=0}^{n} \mu^{e,k}_j q^{e,k}_j\big)^2 + \sum_{p \in P_{k,f}} \big(p - \sum_{j=0}^{n} \mu^{k,f}_j q^{k,f}_j\big)^2 + \sum_{f \in \textit{rhs}(e,k)} d(p_f, p_e + r_k)^2 + \sum_{e \in \textit{rhs}(k,f)} d(p_e, p_f - r_k)^2$$

where we write e.g. $p \in P_{e,k}$ to sum over all entities e and all points p in $P_{e,k}$. We again use nuclear norm regularization to enforce low-dimensional subspaces. Let the i^{th} row vector of the matrix $M_{e,k}$ be given by $q^{e,k}_i - q^{e,k}_0$ and similar for $M_{k,f}$. We define:

$$J^2_{\textit{reg}} = \sum_{k \in R} \|M_{e,k}\|_* + \|M_{k,f}\|_*$$

Note that we only need to consider the combination (e,k) or the combination (k,f) if there is at least one triple of the form (e,k,f) in G, since otherwise we can trivially choose $M_{e,k}$ and $M_{e,k}$ as the zero matrix. The full regularization term is given by $J_{\textit{reg}} = J^1_{\textit{reg}} + J^2_{\textit{reg}}$.

4 Evaluation

In our experiments, we have used Wikidata[5] to obtain a set of entities E and their corresponding semantic types. To generate the bag-of-words representation W_e of a given entity, we

[5]Most knowledge graph embedding models have been evaluated on fragments of Freebase and WordNet. Our choice of Wikidata is motivated by its relatively clean semantic type information. For example, while Barack Obama is of type *Human* on Wikidata, Freebase among others mentions the following types: *film subject*, *musical artist* and *building occupant*. WordNet has a rich semantic type hierarchy, but contains relatively few instances of these types (e.g. of the 51K leaf nodes in WordNet 1.7 only 7K were found to be instances in [Alfonseca and Manandhar, 2002]).

take advantage of the fact that Wikidata entities e are linked to their corresponding Wikipedia article d_e. The set W_e contains the words occurring in d_e, as well as the m words before and after any mentions of the entity in other Wikipedia articles. Following [Pennington *et al.*, 2014], we have used a window size of $m = 10$ (but without crossing sentence boundaries). In particular, we treat every link from some Wikipedia article d_x to d_e as a mention of e, as well as any repeated occurrences of the corresponding anchor text in d_x. The word-word co-occurrence in the J_{glove} component of our model has been obtained from the entire Wikipedia corpus, as in the standard GloVe model. Using the Wikidata dump from October 26, 2015 and the Wikipedia dump from August 5, 2015, we have then selected those entities e which are mentioned in at least 10 Wikipedia articles, resulting in a set E containing 1,292,702 entities. For each semantic type s, the set E_s contains those entities which are asserted to be of type s via the *instance of* property as well as all instances which are asserted to belong to one of the supertypes of s, which are determined using the *subclass of* property. As the set of binary relations R we considered all Wikidata properties whose value is another entity, apart from *instance of* and *subclass of* which have already been used to determine the sets E_s. In the case of Wikipedia, we adopted a fairly straightforward preprocessing strategy, as used in many other works such as [Wang *et al.*, 2014]. In particular, we removed punctuations, lower-cased the tokens, and conducted sentence segmentation using the NLTK library[6]. We also removed words whose term frequency in the entire collection was less than 10.

Our main baseline is pTransE, which also learns an embedding of entities by combining a word embedding model with a knowledge graph embedding model. We consider three variants of this baseline: pTransE$_{\text{anch}}$ is the version proposed in [Wang *et al.*, 2014], which uses anchor text for aligning word vectors and entity vectors; pTransE$_{\text{art}}$ is the improvement proposed in [Zhong *et al.*, 2015], which uses the words in the Wikipedia article d_e instead of anchor text (and a slightly different model); pTransE$_{\text{full}}$ is a variant of pTransE$_{\text{art}}$, which uses the bag of words representation W_e instead, as in our method. In addition, we also compare our method against a number of knowledge graph embedding methods: TransE, TransH, TransR and CTransR. We used Bernoulli sampling for selecting negative examples (see [Wang *et al.*, 2014]); we also evaluated uniform sampling (not shown), and found the results to be very similar to Bernoulli sampling but slightly worse. It is expected that these methods will perform worse, as they cannot exploit the text representation W_e of the entities. We also compare our method with Skip-gram and CBOW, which can only use text representations and are thus also expected to perform worse. In particular, to apply these models to learn entity embeddings, we use the same method as for our model to determine entity mentions on Wikipedia, and then apply the standard models based on the words surrounding these mentions. Finally, we have compared our method with the multi-dimensional scaling (MDS) based method from [Derrac and Schockaert, 2015], in which case we learn a separate vector space for every semantic type.

[6]http://www.nltk.org/

	Ranking	Induction			Analogy
	ρ	MAP	P@5	MRR	Acc.
Skip-gram	0.102	0.176	0.356	0.505	0.184
CBOW	0.105	0.182	0.350	0.500	0.213
TransE	0.110	0.060	0.200	0.451	0.382
TransH	0.142	0.072	0.210	0.415	0.382
TransR	0.100	0.102	0.302	0.489	0.378
CTransR	0.122	0.132	0.323	0.499	0.402
pTransE$_{\text{anch}}$	0.099	0.101	0.301	0.488	0.476
pTransE$_{\text{art}}$	0.202	0.218	0.475	0.751	0.512
pTransE$_{\text{full}}$	0.213	0.224	0.490	0.756	0.532
EECS$_{\text{full}}$	**0.319**	**0.231**	**0.609**	**0.883**	**0.591**
EECS$_{\text{no rel}}$	0.301	0.201	0.586	0.843	0.552
EECS$_{\text{no type}}$	0.266	0.197	0.579	0.801	0.532
EECS$_{\text{no NN}}$	0.220	0.176	0.554	0.789	0.500
EECS$_{\text{no dist}}$	0.309	0.201	0.568	0.821	0.542
EECS$_{\text{text}}$	0.254	0.228	0.579	0.831	0.540

Table 1: Experimental results for the full test data.

Because of the limited scalability of the latter model, however, we have only considered this for semantic types with up to 10000 instances.

To analyze the relative importance of the different components, we consider a number of variants of our model, which we will refer to as EECS (Entity Embeddings with Conceptual Subspaces). EECS$_{\text{full}}$ refers to our full model; EECS$_{\text{no rel}}$ refers to a variant in which J_{rel} and the associated regularization component J^2_{reg} have been removed; EECS$_{\text{no type}}$ refers to a variant in which J_{type} and J^1_{reg} have been removed; EECS$_{\text{no NN}}$ refers to a variant in which the regularization component J_{reg} has been removed (which also trivializes the component J_{type}); EECS$_{\text{no dist}}$ refers to a variant in which the distance constraints of the form $d(p_f, p_e + r_k)^2$ and $d(p_e, p_f - r_k)^2$ in J_{rel} have been removed; finally EECS$_{\text{text}}$ refers to a variant in which only the component J_{text} is used, reducing our model essentially to a variant of GloVe. The implementation and datasets will be made available online.

All experiments were evaluated using five-fold cross validation. For tuning the parameter β of our model, based on a tuning/validation set in each experiment, we considered the range $\{50, 100, 150, 200, 250, 300, 350, 400\}$. For the parameter α, we considered values between 0 and 1 with an increment of 0.1. The number of iterations for all models was set to 20, as we found that beyond this number empirical results became fairly consistent in all cases. Based on the tuning set, in each of the experiments the optimal value of β was found to be 300, while the optimal values of α varied between 0.4, 0.5 and 0.6. The number of dimensions was always set to 300 for our model, noting that because of the nuclear norm regularization this only represents an upper bound on the actual number of dimensions. All parameters of the baseline methods, including the number of dimensions, have been optimized based on the tuning set in each experiment.

Ranking A characteristic feature of conceptual spaces is that they are encoded as Cartesian products of interpretable dimensions. For a vector space model to be meaningful as a conceptual space, it is therefore important that salient proper-

	Ranking	Induction			Analogy
	ρ	MAP	P@5	MRR	Acc.
MDS	0.022	0.054	0.133	0.284	0.201
$\text{EECS}_{\text{full}}$	0.218	0.140	0.301	0.463	0.456

Table 2: Comparison with MDS on a subset of the test data.

ties can be modelled as directions[7]. Therefore, we have evaluated the ability of our model to correctly rank entities according to a given property. As we need the ground truth, we have focused on properties with numerical values which are available in Wikidata (but have not been considered when learning the space), e.g. the *date of birth* for entities of type *human*, or the *boiling point* of entities of type *chemical element*. In total, we have retrieved 26 numerical attributes which are available for at least 30 entities. Some of these numerical attributes appear for different semantic types (e.g. the property *inception* applies to the semantic types *film*, *organization* and *country*, among others). In total, we obtained 73 such property-type combinations, each of which is considered as a problem instance. For each problem instance, the corresponding set of entities is split into 60% training, 20% validation and 20% testing sets. From the training set, a direction is estimated using the SVMRank model[8] [Joachims, 2002]. The parameters of the resulting ranking model are optimized using the validation sets. Table 1 shows the performance on the testing set, in terms of Spearman's ρ[9], expressing the correlation between the ranking predicted by the model and the ranking according to the numerical values found in Wikidata.

The results show that standard word and knowledge graph embedding models are not competitive, which is not surprising given that they use less information than our model. However, the results also show that our model substantially outperforms pTransE, even the variant $\text{pTransE}_{\text{full}}$ which uses the same input as our model. Comparing the results for $\text{EECS}_{\text{full}}$ and $\text{EECS}_{\text{no NN}}$ clearly shows that the nuclear norm regularization is largely responsible for the good performance of our model; indeed the performance of $\text{EECS}_{\text{no NN}}$ is comparable to $\text{pTransE}_{\text{full}}$. Also note that disregarding relations ($\text{EECS}_{\text{no rel}}$) only leads to a small drop in performance. The results in Table 2 compare our model against the MDS model from [Derrac and Schockaert, 2015] on a reduced set of 27 problem instances. These results clearly show that the MDS method is not competitive.

Induction A second characteristic feature of conceptual spaces is that properties correspond to convex regions. Moreover, it is often assumed that the boundaries of these regions are determined based on the distance to a particular point in the space, which acts as a prototype. In this experiment, we test our method's ability to make inductive inferences based on this view. In particular, given a number of entities of the same type which have some (unknown) property in common, the task we consider is to identify other entities that also have this property. Problem instances in this case were obtained by omitting all triples of the form $(., r, f)$ for particular choices of r and f, when learning the embeddings. The set of entities e for which $(e, r, f) \in G$ then defines a problem instance. This set of entities is again split into a training, tuning and testing set. For evaluation purposes, we consider this task as a ranking task. In particular, for each problem instance we rank the entities of the considered type, based on their distance to the center-of-gravity of the training instances, and evaluate the quality of this ranking using mean average precision (MAP), Precision@5 (P@5) and Mean Reciprocal Rank (MRR); note that in all cases, higher values are better.

The results in Table 1 show that $\text{EECS}_{\text{full}}$ again outperforms all of the baselines. Note, however, that in the case of MAP, the differences with $\text{EECS}_{\text{text}}$ and $\text{pTransE}_{\text{full}}$ are relatively small. The fact that the differences are clearer for P@5 and MRR suggests that our method is better able to select a few entities with high precision. The MAP score tends to be dominated by outliers, leading to smaller differences. The MDS model is again not competitive.

Analogy Finally we have considered the problem of completing analogical proportions of the form "a is to b what c is to ...", which is a standard evaluation task for word embeddings. Our main aim in this task is to evaluate how well different subspaces are aligned. We have used the set of problem instances from the GloVe project[10]. However, since we had to restrict ourselves to analogical proportions involving entities, only a subset of 8363 problem instances could be considered. In this case we randomly split the data into 25% tuning and 75% testing sets.

The results again reveal that $\text{EECS}_{\text{full}}$ outperforms all baselines. Note that the performance of $\text{EECS}_{\text{no type}}$ is identical to $\text{pTransE}_{\text{full}}$, which shows that the improvement is mostly due to the the fact that we model semantic types as subspaces. The results for the MDS method are again not competitive.

5 Conclusions

We have proposed a new method for learning vector-space embeddings of entities, based on available semantic information (from Wikidata) and textual descriptions (from Wikipedia). From a technical point of view, the main novelty of our model is the use of nuclear norm regularization to encode the requirement that entities of the same semantic type should lie in a low-dimensional subspace. In particular, nuclear norm regularization allows the model to automatically select the most appropriate number of dimensions for the subspace corresponding to each type. From an application point of view, our main motivation was to learn subspaces that are useful as approximations of conceptual spaces. To support this view, among others, we have shown that many numerical attributes can be faithfully modelled as directions and that the learned representations allow us to model induction based on distance to a centroid.

[7] Conceptual space representations also encode information about the correlation between the underlying dimensions, which in our case is captured by the angles between these directions.

[8] https://www.cs.cornell.edu/people/tj/svm_light/svm_rank.html

[9] The reported average ρ values have been obtained using the Fisher z-transformation.

[10] http://nlp.stanford.edu/projects/glove/

References

[Alfonseca and Manandhar, 2002] Enrique Alfonseca and Suresh Manandhar. Distinguishing concepts and instances in wordnet. In *Proceedings of the First International Conference of the Global WordNet Association*, 2002.

[Bordes *et al.*, 2011] A. Bordes, J. Weston, R. Collobert, and Y. Bengio. Learning structured embeddings of knowledge bases. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011.

[Bordes *et al.*, 2013] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko. Translating embeddings for modeling multi-relational data. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, pages 2787–2795. 2013.

[Chella, 2015] Antonio Chella. A cognitive architecture for music perception exploiting conceptual spaces. In Frank Zenker and Peter Grdenfors, editors, *Applications of Conceptual Spaces*, pages 187–203. Springer International Publishing, 2015.

[Collins and Michalski, 1989] A. Collins and R. Michalski. The logic of plausible reasoning: A core theory. *Cognitive Science*, 13(1):1–49, 1989.

[Derrac and Schockaert, 2015] J. Derrac and S. Schockaert. Inducing semantic relations from conceptual spaces: a data-driven approach to plausible reasoning. *Artificial Intelligence*, pages 74–105, 2015.

[Dong *et al.*, 2014] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmann, S. Sun, and W. Zhang. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 601–610, 2014.

[Douven *et al.*, 2013] I. Douven, L. Decock, R. Dietz, and P. Égré. Vagueness: A conceptual spaces approach. *Journal of Philosophical Logic*, 42:137–160, 2013.

[Fazel, 2002] Maryam Fazel. *Matrix rank minimization with applications*. PhD thesis, Stanford University, 2002.

[Festinger, 1957] Leon Festinger. *A theory of cognitive dissonance*. Stanford University Press, 1957.

[Forth *et al.*, 2010] J. Forth, G. A Wiggins, and A. McLean. Unifying conceptual spaces: Concept formation in musical creative systems. *Minds and Machines*, 20:503–532, 2010.

[Gärdenfors, 2000] P. Gärdenfors. *Conceptual Spaces: The Geometry of Thought*. MIT Press, 2000.

[Guo *et al.*, 2015] S. Guo, Q. Wang, B. Wang, L. Wang, and L. Guo. Semantically smooth knowledge graph embedding. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing*, pages 84–94, 2015.

[Hsieh and Olsen, 2014] Cho-Jui Hsieh and Peder Olsen. Nuclear norm minimization via active subspace selection. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 575–583, 2014.

[Joachims, 2002] Thorsten Joachims. Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 133–142. ACM, 2002.

[Lieto *et al.*, 2015] Antonio Lieto, Andrea Minieri, Alberto Piana, and Daniele P. Radicioni. A knowledge-based system for prototypical reasoning. *Connection Science*, 27:137–152, 2015.

[Lin *et al.*, 2015] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. Learning entity and relation embeddings for knowledge graph completion. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 2181–2187, 2015.

[Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *Proceedings of the 27th Annual Conference on Neural Information Processing Systems*, pages 3111–3119, 2013.

[Pennington *et al.*, 2014] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543, 2014.

[Recht *et al.*, 2010] Benjamin Recht, Maryam Fazel, and Pablo A Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52:471–501, 2010.

[Rosch, 1973] Eleanor H Rosch. Natural categories. *Cognitive Psychology*, 4(3):328–350, 1973.

[Schockaert and Prade, 2013] Steven Schockaert and Henri Prade. Interpolative and extrapolative reasoning in propositional theories using qualitative knowledge about conceptual spaces. *Artificial Intelligence*, 202:86–131, 2013.

[Turney and Pantel, 2010] P. D. Turney and P. Pantel. From frequency to meaning: Vector space models of semantics. *Journal of Artificial Intelligence Research*, 37:141–188, 2010.

[Wang *et al.*, 2014] Z. Wang, J. Zhang, J. Feng, and Z. Chen. Knowledge graph and text jointly embedding. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1591–1601, 2014.

[Xu *et al.*, 2014] C. Xu, Y. Bai, J. Bian, B. Gao, G. Wang, X. Liu, and T.-Y. Liu. RC-NET: A general framework for incorporating knowledge into word representations. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 1219–1228, 2014.

[Zhong *et al.*, 2015] Huaping Zhong, Jianwen Zhang, Zhen Wang, Hai Wan, and Zheng Chen. Aligning knowledge and text embeddings by entity descriptions. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 267–272, 2015.